Randomized Admission Policy for Efficient Top-k and Frequency Estimation

Ran Ben Basat
Computer Science
Technion
sran@cs.technion.ac.il

Gil Einziger
Electrical Engineering
Politechnico di Torino
gilga@polito.it

Roy Friedman
Computer Science
Technion
roy@cs.technion.ac.il

Yaron Kassner
Computer Science
Technion
kassnery@cs.technion.ac.il

arXiv:1612.02962v1 [cs.DS] 9 Dec 2016

***Abstract*—Network management protocols often require timely and meaningful insight about per flow network traffic. This paper introduces *Randomized Admission Policy (RAP)* – a novel algorithm for the *frequency* and *top-k* estimation problems, which are fundamental in network monitoring. We demonstrate space reductions compared to the alternatives by a factor of up to 32 on real packet traces and up to 128 on heavy-tailed workloads. For top-k identification, RAP exhibits memory savings by a factor of between 4 and 64 depending on the workloads' skewness. These empirical results are backed by formal analysis, indicating the asymptotic space improvement of our probabilistic admission approach. Additionally, we present d*-Way RAP*, a hardware friendly variant of RAP that empirically maintains its space and accuracy benefits.**

I. Introduction

A. Background

Network management and traffic engineering protocols rely on flow counters based network monitoring. Examples include effective routing, load balancing, QoS enforcement, network caching, anomaly detection and intrusion detection [1]–[6]. Typically, monitoring utilities track millions of flows [7], [8], and the counter of a monitored flow is updated on the arrival of each of its packets. Often, the most frequently appearing flows, known as *heavy hitters*, are also the most interesting, since their impact on the above is the most crucial.

Maintaining such counters is a challenging task with today's storage technology. The difficulty arises as DRAM is too slow to keep up with line rates, while the faster SRAM is expensive and thus too small for keeping an exact counter for each flow. These limitation were tackled using various approaches.

Estimators reduce the size of counters using probabilistic techniques [9]–[11]. This enables maintaining one counter per flow in SRAM at the cost of reduced accuracy. The downside of estimators is that they require an explicit flow to counter mapping for every flow. This mapping often becomes the dominant factor in memory consumption [12].

The *shared counters* approach, also known as *sketches*, solves the mapping problem using hashing algorithms that implicitly assign flows to counters. Well known examples include *Multi Stage Filters* [13] and *Count Min Sketch* [14]. Yet, to reduce the impact of hash collisions on counters' reading accuracy, these methods must allocate considerably more space and more counters than predicted by lower bounds.

Databases and data analytics face similar problems, known in these domains as *frequency estimation* and *top-k identification*, i.e., identifying who are the k most frequent flows. These domains typically favor *counter based* solutions over sketches since the former are considered superior to sketches, both asymptotically and in practice [15], [16]. Counter based algorithms maintain a fixed size set of counters and aspire to allocate these counters only to the more frequent flows. These include Lossy Counting [17], Frequent [18] and Space Saving [19]. The latter is also considered state of the art [15], [16], [20]. Alas, these algorithms cannot be easily ported into networking devices as they utilize complex data structures and dynamic memory allocation.

Another significant shortcoming of counter based solutions is that they update the state of allocated counters on the arrival of each packet belonging to a unmonitored flow, regardless of how frequent this flow is. Doing so hurts their space to accuracy tradeoff to the point that they become ineffective on heavy-tailed workloads, which are common in network switches and routers.

B. Contributions

In this work, we promote the concept of using a randomized admission policy for allocating counters to non-monitored flows, and show that it can significantly improve accuracy. Intuitively, such a policy ignores most of the tail flows and is still able to eventually admit the high frequency flows.

Specifically, this idea is realized in a novel counter based algorithm called *Randomized Admission Policy* (RAP) as well as a hardware friendly variant called *d-Way associative RAP* (dW-RAP). RAP is simpler to analyze, while dW-RAP maps well into limited associativity cache designs and empirically maintains most of the benefits of RAP. We extensively evaluate RAP and dW-RAP over two real packet traces [21], [22], a YouTube access trace [23] and synthetic Zipf distributions.

For the frequency estimation problem, RAP and dW-RAP achieve the same *mean square error* (MSE) as the leading alternatives while using a fraction of the required memory. For top-k identification, RAP and dW-RAP exhibit significantly higher recall and precision, even when allocated with half the space given to the alternative methods. In particular, when the distribution is only mildly skewed (or heavy-tailed), RAP and

dW-RAP are the only techniques that successfully identify a high percentage of the top-k flows.

II. Related Work

The frequent items and top-k identification problems appear in slight variations across multiple domains. Algorithms for these problems are often categorized as either *counter based* or *sketch based*. In addition, the specific challenges of network monitoring have spawned solutions that are especially tailored for the memory limitations in the networking case.

A. Counter based algorithms

Counter based algorithms are usually designed for software implementations and maintain a table of monitored items. The differences between these algorithms lie in the question of admission and eviction of entries to and from the table. From a networking perspective, counter based algorithms maintain an explicit flow to counter mapping for monitored items. For a stream of N events and an accuracy parameter ε, the goal is to approximate a given flow's frequency to within an additive error of $N \cdot \varepsilon$. For this task, $\Omega(\frac{1}{\varepsilon})$ counters are required [19], and this is achieved by some of the algorithms below.

Lossy Counting [17] increments an arriving item's counter on every arrival. If the counter is not in the table, it is admitted with a counter value of 1. Lossy Counting keeps the table size bounded by periodically decrementing table counters and evicting items whose counter reaches 0. Unfortunately, Lossy Counting requires a maximal number of $\frac{1}{\varepsilon} \cdot \log(N)$ table entries. *Probabilistic Lossy Counting* [24] requires fewer table entries on average but only provides a probabilistic guarantee.

In *Frequent (FR)* [18], [25], whenever an item arrives and the table already contains $\frac{1}{\varepsilon}$ entries, the item is not admitted. Instead, FR decrements every entry in the table, evicting entries whose counter reached 0. The main benefit of FR is that it requires the optimal number of $O(\frac{1}{\varepsilon})$ table entries.

Space Saving (SS) [19] requires the same number of entries as FR, but maintains additional information to improve accuracy. Space Saving admits *any* arriving item at the expense of evicting the minimum-frequency item. Space Saving is considered to be state of the art [15], [16], [20].

B. Sketch based algorithms

Sketches, such as *Multi Stage Filters* [26], *Count Sketch* [27] and *Count Min Sketch* [14], are very common in networking domains as they are simple to implement in hardware and have low implementation overheads. The most popular example, Count Min Sketch, provides the following guarantee — given an item x, with probability of at least $1 - \delta$, the estimation error of x is at most $N \cdot \epsilon$.

Count Min Sketch does not require storing flow identifiers or maintaining a flow to counter association. Instead, it maintains an array of $\ln(\frac{1}{\delta})$ rows, each with $\frac{e}{\varepsilon}$ counters. When an item arrives, a hash function is calculated for each row and its corresponding counter is incremented. To estimate the frequency of an item, the corresponding counters are read and the minimum counter value is returned as the estimation.

Asymptotically, Count Min Sketch requires a suboptimal number of counters. However, it does not store flow ids and has only minor overheads for hardware implementation. Despite being suboptimal, sketches still require a sub-linear number of counters, can completely reside in SRAM, and provide online frequency estimation.

On the contrary, *Counter Braids [12]* and *Counter Tree* [28] use an hierarchical sketch where overflowing counters are hashed to a higher level sketch. They are able to encode items just like Count Min Sketch would, but the decoding process is complex and can only be performed offline, estimating all flow values together.

In *Randomized Counter Sharing [29]*, every time an item is added, a random hash function is used and the corresponding counter is incremented. The flow identifier is recorded, but without an explicit mapping to frequency. When a measurement ends, we estimate the flow's frequency by summing all of the corresponding counters or by performing a maximum-likelihood estimation. Both of these estimations are quite slow and cannot be performed online.

In summary, sketches are space suboptimal and only solve the frequency estimation problem. Further, they only support point queries and their answers are only correct within a certain probability. Despite these limitations, sketches are used for many networking applications [14], [24], [26], [30]–[33].

C. Network monitoring architectures

In *hybrid SRAM/DRAM architectures* [7], [8], the LSB bits of counters are stored in SRAM and the MSB in DRAM. This way, the space allocated for each flow in SRAM is small. However, the SRAM counters have to periodically be synchronized with the DRAM counters, which increases the contention on the memory bus. Further, estimating a flow's frequency requires accessing DRAM and therefore cannot be used for online network monitoring.

Brick [34] uses an efficient encoding in order to reduce the number of bits allocated per counter. Brick enables storing more counters, under the assumption that the total value of increments is known in advance. Brick is most effective when there are many very small flows.

Estimators use fixed size small counters in order to represent large numbers. These methods trade precision for space and allow more counters to be contained in SRAM. This idea was first introduced by *Approximate Counting* [35] and was adapted to networking devices [9]–[11], [36]. The downside of estimators is that they require storing a flow-to-counter mapping for every flow, a requirement that has many overheads. Sampling techniques are another alternative that trades accuracy for space. Unfortunately, these methods can only monitor large flows that are frequent enough to be sampled [37], [38].

III. Randomized Admission Policy (RAP)

RAP maintains a table (C) which contains $M \triangleq \frac{1}{\epsilon}$ entries. The intuition behind RAP is to minimize the error inflicted upon arrival of a non-monitored item $x \notin C$ when the table is full. That is, we identify inefficiencies in the way previous

works behave in this case. E.g., FR needlessly increases the error of all counters by decrementing all of them. In contrast, Space Saving always evicts the item with the minimal counter. This eviction introduces an error, as the monitored element is often more frequent than a randomly arriving item without a counter. This is especially true for heavy-tailed workloads, where a large fraction of the stream consists of "tail elements" that should not be admitted into the table. In RAP, we take a more conservative approach. When an item $x \notin C$ arrives, we find the item (m) with the minimal counter value (c_m). x is then admitted into C with probability $\frac{1}{c_m+1}$ at the expense of m; otherwise, x is simply discarded. Algorithm 1 provides a pseudo code of the RAP's ADD method.

In order for an item x to replace the minimal element m, it has to arrive $c_m + 1$ times on average. Infrequent items are therefore unlikely to be admitted into C, and most of them will not affect any of the counters. Therefore, RAP is considerably more accurate, especially for heavy tailed workloads where a large portion of the items are infrequent. In contrast, every tail item in Space Saving affects the counters, thereby contributing to the total estimation error. Our approach is not without risks, as if an arriving item turns out to be frequent, Space Saving admits that item sooner than RAP.

Given a query for the frequency of element x, RAP estimates it as c_x if $x \in C$ and 0 otherwise.

RAP can be implemented with existing data structures and it processes packets at $O(1)$ runtime [19], [32]. It stores a single counter per table entry, while Space Saving entries are slightly larger as they store two values.

Algorithm 1 Randomized Admission Policy

Initialization: $C \leftarrow \emptyset, \forall i : c_i \leftarrow 0$

function ADD(Item x)
  if $x \in C$ **then**
    $c_x \leftarrow c_x + 1$
  else
    if $|C| < M$ **then**
      $c_x \leftarrow 1$
      $C \leftarrow C \cup \{x\}$
    else
      $m \leftarrow \text{argmin}_{y \in C} c_y$
      if $random() < \frac{1}{c_m+1}$ **then** ▷ w.p $\frac{1}{c_m+1}$.
        $C \leftarrow (C \setminus \{m\}) \cup \{x\}$
        $c_x \leftarrow c_m + 1$
      end if
    end if
  end if
end function

A. Analysis

We start our analysis with theoretic bounds for the top-k problem. These show that our probabilistic approach is asymptotically better for i.i.d. streams. We then explore the properties of RAP for the frequency estimation problem.

1) Top-k Problem: We say that an algorithm successfully solves top-k if it identifies the k most frequent flows in a stream. Our goal is to bound the number of table entries required for successful identification of top-k.

Denote $\Gamma_\alpha(D) \triangleq \sum_{i=1}^{D} i^{-\alpha}$. A stream will be called an i.i.d. Zipf stream with skew α over domain D if all of its elements are sampled independently and follow the distribution in which item $i \in \{1, 2, \ldots, D\}$ appears with probability $\frac{i^{-\alpha}}{\Gamma_\alpha(D)}$. Such a stream is denoted $\mathcal{Z}_\alpha^D$.

Theorem III.1. *Let $\alpha < 1$. For any fixed k, Space Saving requires $O(D^{1-\alpha})$ counters to solve top-k on $\mathcal{Z}_\alpha^D$, while a randomized admission policy requires only to $O(D^{\frac{1-\alpha}{1+\alpha}})$.*

Next, we analyze the performance of a probabilistic admission filter for streams with higher skew. The proofs of both theorems appear in Appendix A.

Theorem III.2. *For $\mathcal{Z}_1^D$ and any fixed k, Space Saving requires $O(\log D)$ counters to solve top-k, while a randomized admission policy reduces the required number to $O(\sqrt{\log D})$.*

2) Frequency Estimation Problem: We now present a brief mathematical analysis of RAP, including deterministic and probabilistic upper bounds for the estimation error.

Theorem III.3. *Let f_x be the true frequency of x, $\widehat{f_x}$ be RAP's estimation of f_x, and m be $\min_{y \in C} C_y$. Then $\widehat{f_x} \leq f_x + m$.*

Proof. The proof is by a case analysis. First, suppose $x \notin C$ at the time of the query. In this case, $\widehat{f_x} = 0$ and the claim trivially holds. Conversely, assume that $x \in C$ at the time of the query; consider the last time t in which x was admitted into C. At that point, $c_x^t = m_{t^-} + 1$ where c_x^t is the value of x's counter at time t and m_{t^-} is the minimum counter in the table just before time t. Notice that the algorithm can only increase the minimum counter in the table due to a packet arrival, at which point either no counter changes or the minimal counter is incremented. Hence, $c_x^t \leq m_{t^-} + 1$. Next, suppose that x has arrived exactly n times between t and the present; $n \leq f_x - 1$ since we know that at time t, x arrived once. It follows that $c_x = c_x^t + n \leq m_{t^-} + 1 + f_x - 1 = f_x + m_{t^-} \leq f_x + m$. □

Next, we show that the estimation given by RAP is in expectation smaller than or equal to the true frequency.

Theorem III.4. $\mathbb{E}\left[\widehat{f_x}\right] \leq f_x$.

Proof. In this proof, we use the notion of time to describe the events in the stream. The first event is at $t = 0$, the next at $t = 1$ and so on. We prove the claim by induction on the time t. *Base:* At time 0, $\widehat{f_x} = 0$. *Step:* If at time t an item different than x has arrived, then $\mathbb{E}\left[\widehat{f_x^t}\right] \leq \mathbb{E}\left[\widehat{f_x^{t-1}}\right]$; in case c_x was the smallest counter at time $t-1$, its estimation can only decrease, and otherwise its estimation does not change. However, if x arrived at time t, let ΔE_x^t be the change in $\mathbb{E}\left[\widehat{f_x}\right]$, that is $\Delta E_x^t = \mathbb{E}\left[\widehat{f_x^t} - \widehat{f_x^{t-1}}\right]$.

There can be two cases: if $x \in C$, then $\widehat{f_x^t} - \widehat{f_x^{t-1}} = 1$. Otherwise $x \notin C$, hence its estimation is either increased by

$c_m + 1$ with probability $\frac{1}{c_m+1}$ or remains the same. Thus, in all cases $\widehat{f}_x^t - \widehat{f}_x^{t-1}$ grows by 1 in expectation and $\Delta E_x^t = 1$. Hence, the induction hypothesis holds and $\mathbb{E}\left[\widehat{f}_x\right] \leq f_x$. □

IV. Hardware Friendliness

RAP can be efficiently implemented in software with existing data structures [19], [32]. These complex data structures might be difficult to efficiently implement in hardware.

In this section, we present *d-Way Randomized Admission Policy* (dW-RAP), a hardware friendly variant of RAP. We describe dW-RAP as a cache management policy. Caches are well understood, making dW-RAP implementation as a cache policy easy to design as it does not rely on complex data structures. In addition, caches have a proven capability to operate at line speed. For self containment, Section IV-A provides a brief introduction to cache topology.

A. Cache Memory Organization

In order to meet their high speed requirements, hardware caches are usually not fully associative. As a rule of thumb, the higher the associativity level – the slower the cache is since the search process becomes more complex. Limited associativity means that each item can only be placed in a certain logical place in the cache. If this place is already full, an existing item must be evicted in order to admit the new one.

These logical locations are called *sets* and in each set there are a certain number of places called *ways*. We use a hash function ($Set(x)$) to map an item to a certain set number; the item can only be stored in that set. This makes the lookup process simpler as we only need to search for the item in a specific set, rather than in the entire cache.

The more ways we add to the cache – the slower the cache works, as there are more places that an item could be found in. Therefore, to ensure fast performance, the number of ways is kept small, typically $2 - 32$. A cache with d different ways is called *d-way set associative*, a cache with only a single set is called *fully associative*. A cache with a single way is called *direct mapped*.

Figure 1 illustrates the basic topology of a 4-way set associative cache. In this example, the Set function is used to determine the set for x. The set selected is the one marked with orange (horizontal line) and since the cache has 4 ways, x could be placed in either of these ways. The cache first checks whether x appears in these ways. If it is not found, a cache policy is used to decide whether to admit x into the cache, at the expense of evicting some other item, or not.

B. Cache Policy

A fundamental cache management question is what to do when an item arrives and its corresponding set is full. A *cache policy* is an algorithm that answers these questions. Cache policies can sometimes be partitioned into two sub policies: an *admission policy* and an *eviction policy* [6]. The former decides whether to admit an item into the cache and the latter decides on the cache *victim*.

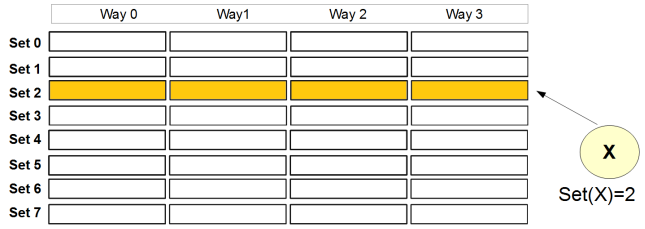

Fig. 1: A 4-way set associative cache with 8 lines. When an item (x) arrives, Set(x)=2 is calculated and the item can be stored in any of the ways of set 2.

C. dW-RAP as a Cache Policy

Algorithm 1 implements RAP assuming (implicitly) a fully associative memory organization. We now describe dW-RAP as a cache policy for a d-way cache organization.

a) Metadata: In dW-RAP, each entry contains a counter that is used for both frequency/top-k estimation, and for the cache admission and eviction policies.

b) Metadata Update: In dW-RAP, every time a cached item is accessed, including right after the initial admission, its counter is incremented by 1.

c) Eviction Policy: When a set is full the cache victim is always the entry with the minimal counter in the set.

d) Admission Policy: dW-RAP's cache policy does not always admit an item into the cache. Instead, it first identifies the set entry with minimal counter as a potential cache victim. If that entry's ID is m and its counter value is c_m, a new item is admitted with probability $\frac{1}{c_m+1}$. The counter of a new item remains with the same value (c_m), and is later incremented by the metadata update.

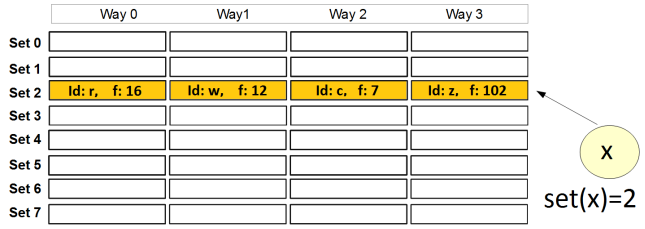

Fig. 2: A 4-way set associative cache with RAP policy. Item x has Set(x)=2. Set 2 includes r, w, c and z, while x is not in the cache. The eviction policy selects c, because its frequency is the smallest. x will be admitted into the cache with probability $\frac{1}{8}$. If x is admitted, its counter will be incremented to 8.

An example of dW-RAP is given in Figure 2. When x arrives, we first look for it in $Set(x) = 2$. There are 4 items in set 2 and x is not one of them. Thus we need to decide on eviction and admission. If we choose to admit x, we evict the minimal item (c) in way 2. The frequency of c is $c_c = 7$ and therefore x is only admitted into the cache with probability $\Pr[Admit(x)] = \frac{1}{c_c+1} = \frac{1}{8}$.
If x is admitted into the cache, c is removed from the cache and c_x is set to be $c_c + 1 = 8$.

Fortunately, the complexity of implementing dW-RAP no longer depends on the number of counters, but only on the associativity level (d). The larger d is, the more combinatoric logic is used for searching the cache and identifying the minimum. As mentioned above, d is typically very small and we can treat the complexity as $O(1)$. In Section V, we evaluate dW-RAP and show that it is almost as accurate as (the fully associative) RAP, even for relatively small values of d. We have also experimented with different associativity levels and evaluated their impact. The experiment details and results appear in Appendix B.

V. Evaluation

In this section, we evaluate RAP and dW-RAP along with the following previously suggested algorithms – Frequent (*FR*) [18] and Space-Saving (*SS*) [19]. These counter-based algorithms were proven effective for both frequency and top-k estimation. The latter is considered state of the art [15], [16], [20]. For frequency estimation, we also compare with sketches such as (*CS*) [27] and (*CMS*) [14].

For a fair comparison, we evaluate the performance of CS and CMS using 8 times as many counters as the rest of the (counter based) algorithms. To represent their low implementation overhead, they were configured to use 4 lines, which was shown effective in practice [39]. By giving the sketches more counters, we compensate for their lower overheads, as they do not maintain a flow to counter association and avoid storing flow identifiers. We consider this a generous comparison, as the flow id and metadata overheads should not take more than 7 times the counter size.

A. Datasets

Our evaluation includes the following datasets:

1) The CAIDA Anonymized Internet Trace 2015 [21] , or in short, *CAIDA*. The data is collected from the 'equinix-chicago' high-speed monitor and contains 18M elements of mixed UDP, TCP and ICMP packets.
2) The UCLA Computer Science department packet trace (denoted *UCLA*) [22]. This trace contains 32M UDP packets passed through the border router of the CS Department, University of California, Los Angeles.
3) YouTube Trace from the UMass campus network (referred to as *YouTube*) [23]. The trace includes a sequence of 600K accesses to YouTube from within the university.
4) Zipf streams. Self-generated traces of identical and independently distributed elements sampled from a Zipf distribution with various skew values ($0.6, 0.8, 1.0, 1.2$ and 1.5). Hereafter, the skew X stream is denoted Zipf$X$.

B. Metrics

Our evaluation considers the following performance metrics:

1) **On-Arrival frequency estimation**
Many networking applications take decisions on a *per-packet* basis. For example, if a router identifies excessive traffic originating from a specific source, the router may suspend further routing of its packets to prevent denial of service attacks. We refer to this as the *On-Arrival* model, where upon arrival of each packet, the algorithms are required to estimate its flow frequency. Formally, a stream $S = s_1, s_2, \ldots$ is revealed one element at a time; consequent to s_t arrival, an algorithm *Alg* is required to provide an estimate $\widehat{f_{s_t}}$ for the number elements in the stream with the same id. We then measure the *Mean Square Error (MSE)* of the algorithm, i.e., $MSE(Alg) \triangleq \frac{1}{N}\sum_{t=1}^{N}(\widehat{f_{s_t}} - f_{s_t})^2$.
2) **Top-k Identification**
The ability to identify the most frequent flows is also important to many applications. We define the *Top-k* identification problem as follows: Given a stream $S = s_1, s_2, \ldots$ and two query parameters m and k, the algorithm is required to output a set of m elements containing as many of the k most frequent stream elements as possible. We denote the k-highest element frequency by F_k. For a set of candidates C, we measure its quality using the standard recall and precision metrics:

$$\begin{aligned} \text{Precision}(C) &\triangleq |e \in C : f_e \geq F_k|/|C| \\ \text{Recall}(C) &\triangleq |e \in C : f_e \geq F_k|/k. \end{aligned}$$

C. On Arrival Evaluation

We begin our evaluation with the frequency estimation problem. In this section, each data point was generated by averaging 10 disjoint batches of 1 million packets each, with the exception of YouTube, which is averaged over two 300,000 batches due to the small size of that trace.

Figure 3 shows the MSE obtained by the different schemes when equipped with an increasing number of counters. We experimented with different associativity levels to conclude that 16W-RAP behaves almost as good as (the fully associative) RAP. The experiment, appearing in Appendix B, shows that while accuracy does improve as associativity grows, there is only little gain by increasing the associativity beyond 16. [1] This remains true under all of our tested workloads.

Figure 9a illustrates our results on the CAIDA packet trace. For the entire range, RAP and 16W-RAP offer significantly lower error than the alternatives. Among the alternatives, none seems to be superior to the rest, as CS, CMS, FR and SS all have some settings where they are more accurate than the rest.

Figure 9d describes the results on the UCLA packet trace. In this trace, RAP is the leader for small memory configurations while for 512 counters and onwards SS is slightly better. We believe this is due to the very high skewness of the trace, meaning that the 512 most frequent elements already consist a significant fraction of the stream and therefore our randomization approach is not needed.

The results for the YouTube traces are illustrated in Figure 9c. In this trace, RAP and 16W-RAP are more accurate than the alternatives. Looking only at the alternatives, it is unclear which is the leading among them. However, they all require more than x4 the space to match the accuracy of RAP.

[1] In modern CPUs, caches have at least 8 ways and 64B lines; thus, each line can accommodate 2 entries, yielding a 16-way structure for our purposes.

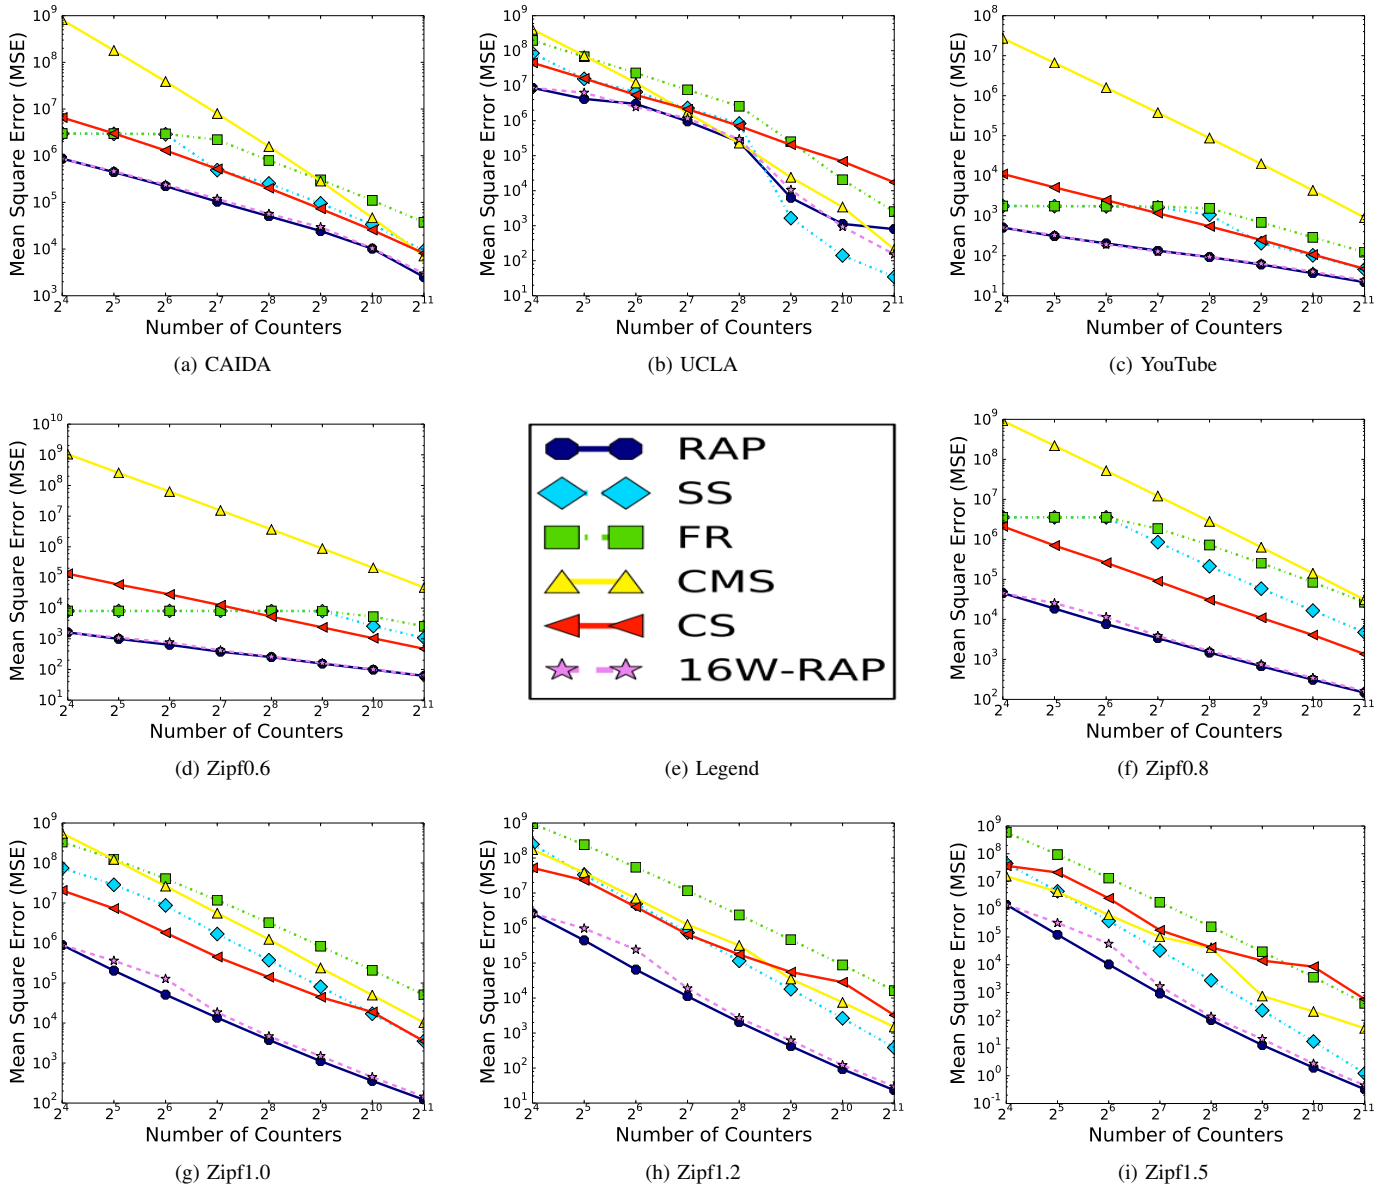

(a) CAIDA (b) UCLA (c) YouTube

(d) Zipf0.6 (e) Legend (f) Zipf0.8

(g) Zipf1.0 (h) Zipf1.2 (i) Zipf1.5

Fig. 3: On Arrival – Mean square error vs. number of counters.
Note that CMS and CS have 8 times as many counters!

e) Synthetic Traces: Synthetic Zipf traces provide us with better insight on the conditions where RAP works best. The least skewed shown distribution is Zipf 0.6 in Figure 9b and the most skewed distribution is Zipf 1.5 in Figure 3i. It appears that RAP performs very well in all these distributions while the alternatives only perform well when the distribution is skewed enough. Figure 9b shows that for Zipf 0.6, SS with 2048 counters obtains worse MSE than RAP with 32 counters! In Figure 3e, we see that 2048 counter SS is about as accurate as a 128 counters RAP. Figure 3g exhibits that for Zipf 1, RAP with 256 counters has similar accuracy as SS with 2048 counters. The trend continues until in Figure 3i the accuracy of RAP with 1024 counters is similar to SS with 2048. For the entire range, RAP requires significantly less space.

D. Top-k Identification

Since, CS and CMS do not solve the top-k problem, we only compare RAP and 16W-RAP to SS and FR. In some of the figures we also gave RAP half the number of counters as the rest. That configuration is marked as 0.5-RAP.

f) Top-32: First, we consider identifying the top-32 flows. We measure the obtained recall for a given number of counters. Our results, summarized in Figure 4, demonstrating that 16W-RAP is almost as accurate as RAP in all workloads.

Figure 4a presents results for CAIDA. As shown, RAP and 16W-RAP achieve near perfect recall with 128 counters. In contrast, FR and SS require 1024 counters for the same recall.

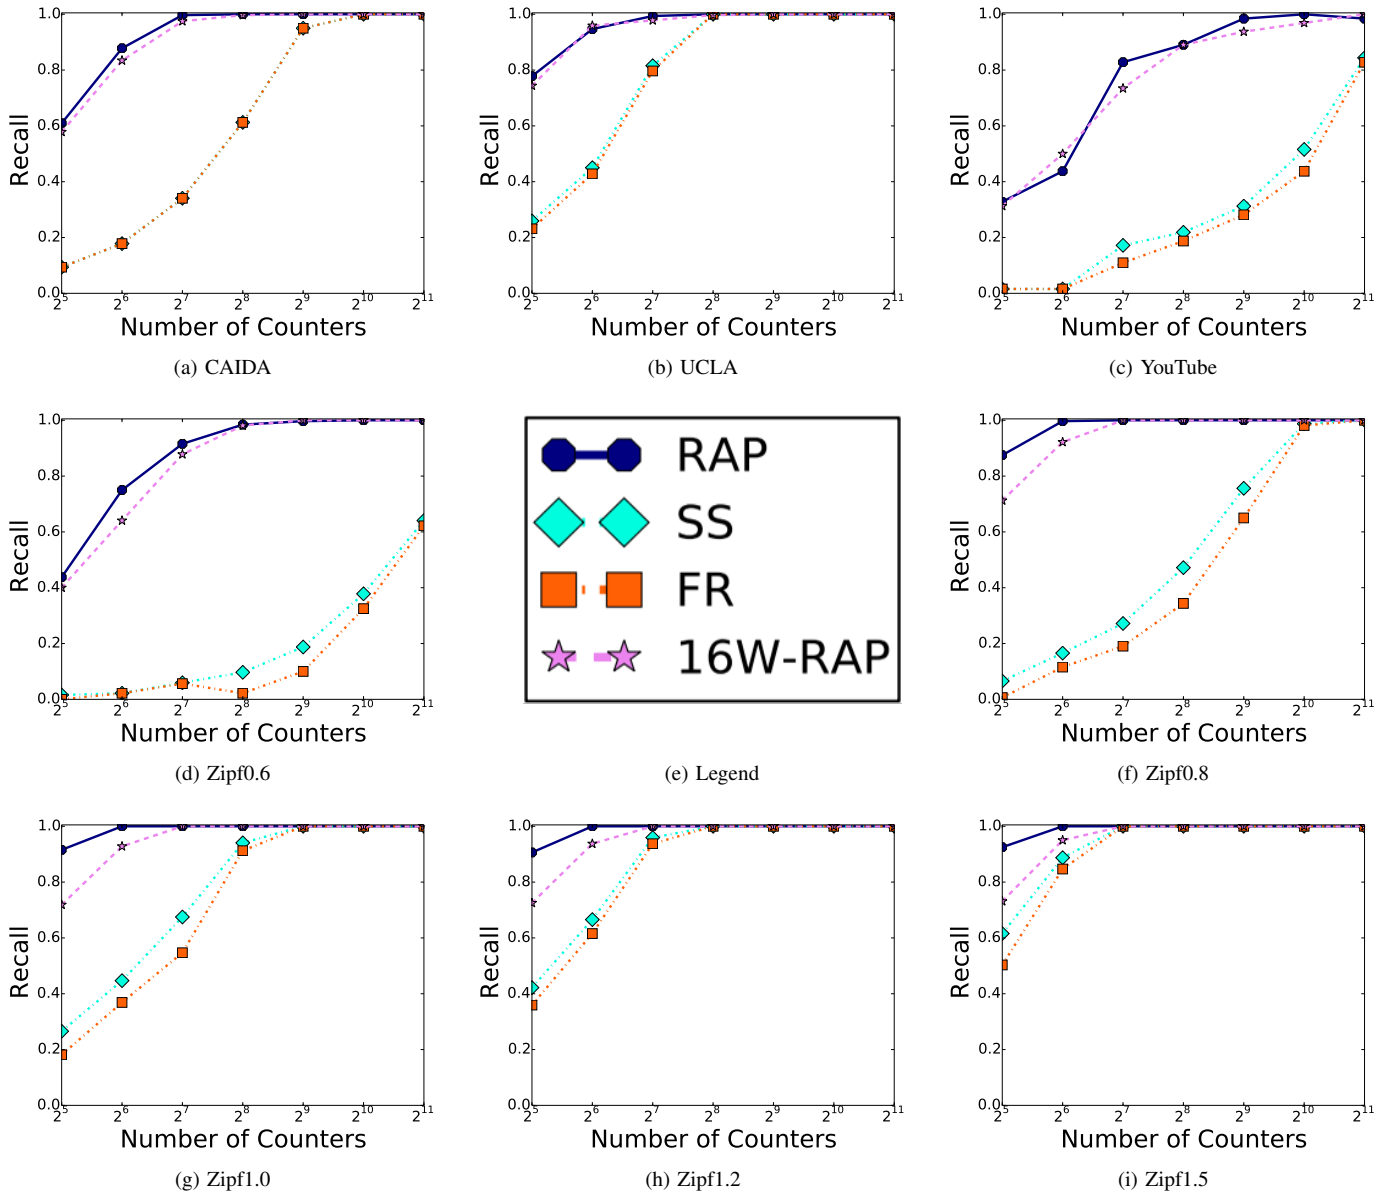

(a) CAIDA (b) UCLA (c) YouTube

(d) Zipf0.6 (e) Legend (f) Zipf0.8

(g) Zipf1.0 (h) Zipf1.2 (i) Zipf1.5

Fig. 4: The recall for finding the top-32 elements vs. number of counters.

Results for UCLA are in Figure 4b. As can be seen, RAP and 16W-RAP reach near optimal recall with 128 counters while FR and SS require 256 counters.

Figure 4c shows that in the YouTube workload, 1024 counters are not enough for SS and FR, and they reach less than 0.5 recall with 1024 counters. In comparison, RAP and 16W-RAP reach near perfect recall with 512 counters.

We now use synthetic Zipf distributions to characterize the performance of RAP and 16W-RAP. Figure 4d shows that for the mildly skewed Zipf 0.6, both RAP and 16W-RAP achieve near optimal recall with 256 counters, while for the alternatives even 2048 are not enough. In the more skewed Zipf 0.8 distribution, Figure 4f shows that RAP requires 64 counters, and 16W-RAP requires 128 to achieve near perfect recall. In contrast, SS and FR require 1024 counters to do the same.

For Zipf 1.0 distribution, Figure 3g shows that RAP and 16W-RAP still require 64 and 128 counters to achieve near optimal recall. FR and SS now require 512 counters to do the same. In Zipf 1.2 distribution, Figure 3h shows that RAP and 16W-RAP continue to require 64 and 128 counters, while SS and FR now require 256 counters to achieve near optimal recall. Finally, for the very skewed Zipf 1.5, Figure 3i shows that RAP and 16W-RAP still require 64 and 128 counters to achieve near optimal recall while SS and FR now require 128 counters. To sum it up, RAP shows a reduction of 2x-16x that depending on workload skewness.

g) Convergence Speed: Since RAP is a randomized algorithm, its convergence speed is as important as its perfor-

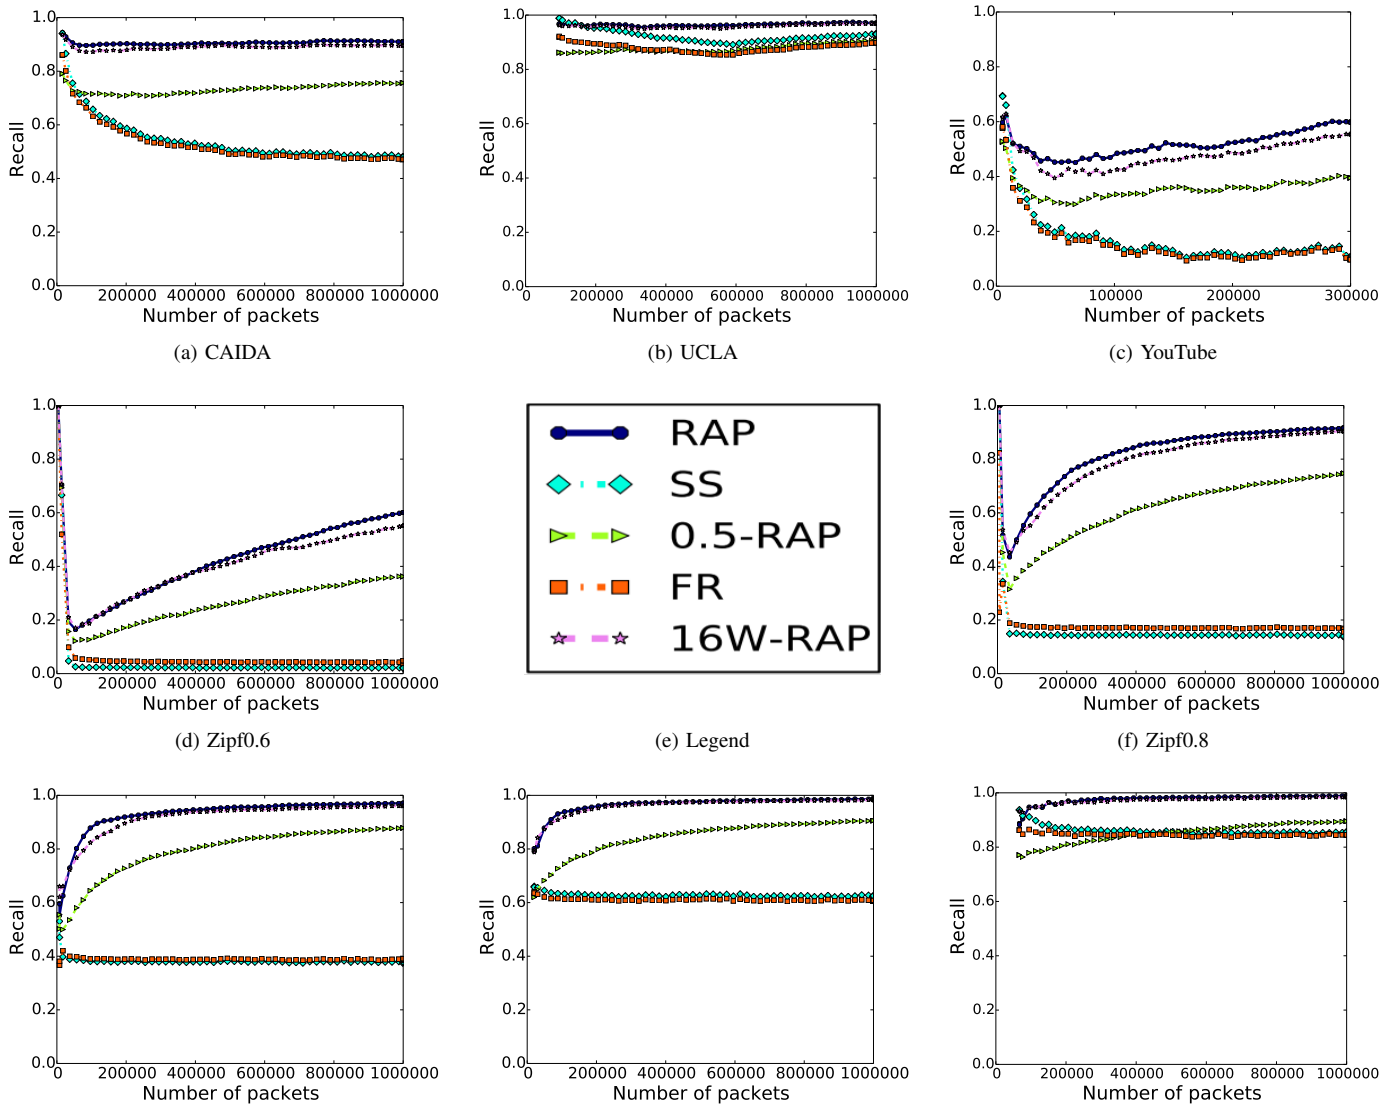

(a) CAIDA (b) UCLA (c) YouTube

(d) Zipf0.6 (e) Legend (f) Zipf0.8

(g) Zipf1.0 (h) Zipf1.2 (i) Zipf1.5

Fig. 5: The average recall achieved for identifying the top-512 using 1024 counters, compared to the size of the stream.

mance for large streams. In order to evaluate the convergence speeds of the different algorithms, we consider the problem of identifying the top-512 flows with 1024 counters. Note that the top-512 items also change during the trace and the algorithms are required to constantly adjust. To quantify the improvement of RAP in these settings we use 0.5-RAP, which attempts to identify the top-512 with mere 512 counters. Our results are illustrated in Figure 5. As can be seen, RAP and 16W-RAP offer similar recall rates in all the tested workloads.

Figure 5a exhibits the results for CAIDA workload. As shown, RAP, 16W-RAP and even 0.5-RAP are significantly better than SS and FR. In Figure 5b, we see that in the UCLA workload, RAP and 16W-RAP achieve around 97% recall while SS and FR achieve less than 92%. Interestingly, in this trace, even 0.5-RAP is above 90% recall with just 512 counters.

Figure 5c shows results for the difficult YouTube trace. Yet, RAP and 16W-RAP achieve above 50% recall, while SS and FR are constantly under 20% recall. The improvement is more than x2, as 0.5 RAP is significantly better than FR and SS.

We now use synthetic workloads to identify the performance envelope of RAP and 16W-RAP. We can observe that as the workload becomes more skewed, all the algorithms improve but RAP and 16W-RAP remain considerably better than the alternatives. For the mildly skewed Zipf 0.6 distribution, Figure 5d shows that RAP and 16W-RAP achieve over 50% recall while FR and SS are under 10%. In the slightly more skewed Zipf 0.8, Figure 5f exhibits that RAP and 16W-RAP yield over 90% recall, while the alternatives are slightly below 20%. Similarly, in Zipf 1.0, Figure 5g shows that RAP and

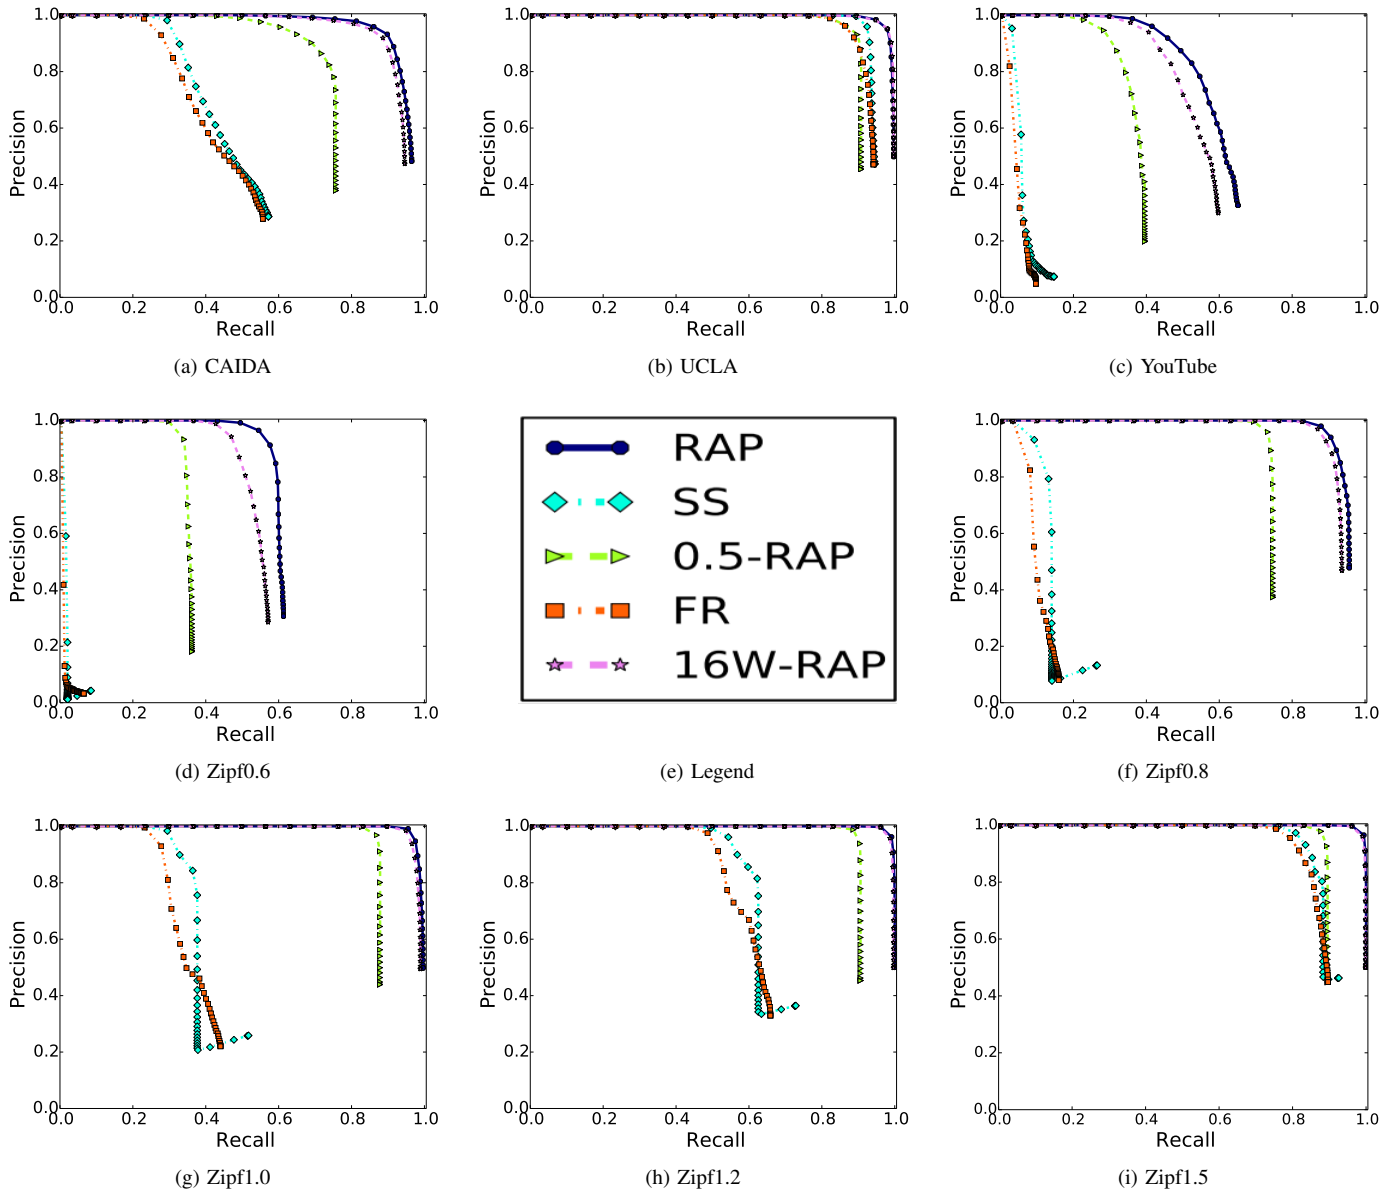

(a) CAIDA (b) UCLA (c) YouTube

(d) Zipf0.6 (e) Legend (f) Zipf0.8

(g) Zipf1.0 (h) Zipf1.2 (i) Zipf1.5

Fig. 6: The Precision-Recall curves for identifying the top-512 with 1024 counters.

16W RAP are close to 100% recall during the entire trace while SS and FR are only at 40% recall. In Zipf 1.2, Figure 5h shows that SS and FR reach 60% while RAP and 16W-RAP remain at 100%. Finally for Zipf 1.5, Figure 5i shows that SS and FR are approaching 85% recall. Even in this workload, 0.5-RAP eventually achieves higher recall. In all these cases, 0.5-RAP achieves higher recall than SS and FR and the space reduction is at least x2.

h) Precision and recall trade-off: While recall is an important measure of the algorithms success, when more than k items are reported as suspected top-k, the precision is compromised. Returning all the items yields the maximum recall, but also poor precision when many items are monitored. The ideal behavior of a top-k algorithm is 100% precision and 100% recall (the top right position in the graphs). Figure 6 illustrates the precision and recall trade-off. For each recall level, we measure how many elements were returned to achieve it, and compute the corresponding precision.

Figures 6a shows our results for the CAIDA workload. RAP and 16W-RAP perform the best on this workload as they can provide 80% recall with near 100% precision, or $\approx 90\%$ recall with $\approx 90\%$ precision. At the same time, their maximum recall is close to 100%, but returning all items drops precision to 50%. SS and FR perform worse as their maximum recall is 60% and they can only ensure 30% recall with high precision.

Figure 6b shows results for the UCLA workload. As can be seen, RAP and 16W-RAP offer the best precision and recall trade-off, although in this case SS and FR also perform well.

Figure 6c shows results for the YouTube workload. As can be observed, this trace is significantly more difficult. RAP and 16W-RAP perform better than the rest. Their maximal recall is slightly over 60% but over 50% recall is possible with very high precision. SS and FR achieve poor recall and precision.

We now look into what happens in synthetic traces. For Zipf 0.6 distribution, Figure 6d shows that RAP and 16W-RAP can achieve over 50% recall with good accuracy, while SS and FR achieve less than 10% recall. In the slightly more skewed Zipf 0.8, Figure 6f shows that RAP and 16W RAP perform very well . They can offer $\approx 90\%$ recall and precision. SS and FR improve slightly, both offer bad accuracy but the maximum recall of SS is 30% and FR is slightly less than 20%. The non-monotone rise in the SS curve is explained by coincidentally having higher rates of top-512 elements in the lower estimated frequency counters. For Zipf 1.0, 1.2 and 1.5, Figure 6g, Figure 6h and Figure 6i show that while RAP and 16W-RAP provides near optimal precision and recall, SS and FR gradually improve as the skew increases. They achieve $\approx 50\%$ recall at Zipf 1.0, $\approx 70\%$ recall at Zipf 1.2 and slightly over 80% recall with high precision with Zipf 1.5. However, in all these cases, 0.5 RAP is better than FR and SS and thus the space reduction is more than x2 across the entire range.

VI. Conclusion and Discussion

In this paper, we have presented Randomized Admission Policy (RAP), a novel algorithm for approximate frequency estimation and top-k identification. We have also introduced d-Way Randomized Admission Policy (dW-RAP), a hardware friendly variant of RAP. We have extensively evaluated RAP and dW-RAP for both problems under two packet traces and a YouTube video trace as well as multiple synthetic Zipf traces. These experiments exhibited significant reductions in the memory requirements of RAP and dW-RAP compared to state of the art alternatives for obtaining the same error. In top-k, we showed that our algorithms achieve superior precision/recall than the alternatives in any tested situation. In the case of frequency estimation, the only exception is the highly skewed UCLA trace [22], and even there, it is only when all schemes are allocated a very large number of counters compared to the trace. Notice that for this case, all algorithms are precise since with many counters on such a skewed trace, the problem becomes almost trivial. In contrast, RAP and dW-RAP are the only algorithms that performed well on heavy-tailed distributions, which are common in Internet services.

Another benefit of RAP and dW-RAP is that they incur fewer updates to memory since they do not replace a counter with each untracked item. This is especially true in heavy-tailed workloads. We have not included the evaluation and quantification of this property for lack of space.

Since dW-RAP can be implemented as a simple cache policy, in the future we would like to integrate it into real networking devices. Interestingly, we believe that dW-RAP may also offer benefits for software implementations. For example, it can probably be parallelized efficiently since each operation only computes the minimum over a small counter set.

References

[1] A. Kabbani, M. Alizadeh, M. Yasuda, R. Pan, and B. Prabhakar, "Af-qcn: Approximate fairness with quantized congestion notification for multi-tenanted data centers," in *IEEE HOTI*, 2010, pp. 58–65.

[2] B. Mukherjee, L. Heberlein, and K. Levitt, "Network intrusion detection," *Network, IEEE*, vol. 8, no. 3, pp. 26–41, 1994.

[3] T. Benson, A. Anand, A. Akella, and M. Zhang, "Microte: Fine grained traffic engineering for data centers," in *ACM CoNEXT*, 2011, p. 8.

[4] P. Garcia-Teodoro, J. E. Daz-Verdejo, G. Maci-Fernndez, and E. Vzquez, "Anomaly-based network intrusion detection: Techniques, systems and challenges," *Computers and Security*, pp. 18–28, 2009.

[5] G. Dittmann and A. Herkersdorf, "Network processor load balancing for high-speed links," in *SPECTS*, vol. 735, 2002.

[6] G. Einziger and R. Friedman, "Tinylfu: A highly efficient cache admission policy," in *Euromicro PDP*, 2014, pp. 146–153.

[7] D. Shah, S. Iyer, B. Prabhakar, and N. McKeown, "Maintaining statistics counters in router line cards." *IEEE Micro*, pp. 76–81, 2002.

[8] S. Ramabhadran and G. Varghese, "Efficient implementation of a statistics counter architecture," *ACM SIGMETRICS*, pp. 261–271, 2003.

[9] G. Einziger, B. Fellman, and Y. Kassner, "Independent counter estimation buckets," in *IEEE INFOCOM*, April 2015, pp. 2560–2568.

[10] E. Tsidon, I. Hanniel, and I. Keslassy, "Estimators also need shared values to grow together," in *IEEE INFOCOM*, 2012, pp. 1889–1897.

[11] C. Hu, B. Liu, H. Zhao, K. Chen, Y. Chen, C. Wu, and Y. Cheng, "Disco: Memory efficient and accurate flow statistics for network measurement," in *IEEE ICDCS*, 2010, pp. 665–674.

[12] Y. Lu, A. Montanari, B. Prabhakar, S. Dharmapurikar, and A. Kabbani, "Counter braids: a novel counter architecture for per-flow measurement," in *ACM SIGMETRICS*, 2008, pp. 1799–1807.

[13] C. Estan and G. Varghese, "New directions in traffic measurement and accounting," *ACM SIGCOMM*, 2002.

[14] G. Cormode and S. Muthukrishnan, "An improved data stream summary: The count-min sketch and its applications," *J. Algorithms*, vol. 55, pp. 29–38, 2004.

[15] G. Cormode and M. Hadjieleftheriou, "Finding frequent items in data streams," *VLDB*, vol. 1, no. 2, pp. 1530–1541, Aug. 2008.

[16] ——, "Methods for finding frequent items in data streams," *J. VLDB*, vol. 19, no. 1, pp. 3–20, 2010.

[17] G. S. Manku and R. Motwani, "Approximate frequency counts over data streams," in *Proc. of the Int. Conf. on V.L. Data Bases*, ser. VLDB, 2002.

[18] R. M. Karp, S. Shenker, and C. H. Papadimitriou, "A simple algorithm for finding frequent elements in streams and bags," *ACM Trans. Database Syst.*, vol. 28, no. 1, Mar. 2003.

[19] A. Metwally, D. Agrawal, and A. E. Abbadi, "Efficient computation of frequent and top-k elements in data streams," in *IN ICDT*, 2005.

[20] N. Manerikar and T. Palpanas, "Frequent items in streaming data: An experimental evaluation of the state-of-the-art," *Data Knowl. Eng.*, pp. 415–430, 2009.

[21] "The caida ucsd anonymized internet traces 2015 - february 19th."

[22] "Unpublished, see http://www.lasr.cs.ucla.edu/ddos/traces/."

[23] M. Zink, K. Suh, Y. Gu, and J. Kurose, "Watch global, cache local: Youtube network traffic at a campus network: measurements and implications," in *Electronic Imaging 2008*. International Society for Optics and Photonics, 2008, pp. 681 805–681 805.

[24] X. Dimitropoulos, P. Hurley, and A. Kind, "Probabilistic lossy counting: An efficient algorithm for finding heavy hitters," *ACM SIGCOMM*, vol. 38, no. 1, Jan. 2008.

[25] E. D. Demaine, A. López-Ortiz, and J. I. Munro, "Frequency estimation of internet packet streams with limited space," in *EATCS ESA*, 2002, pp. 348–360.

[26] C. Estan and G. Varghese, "New directions in traffic measurement and accounting," *ACM SIGCOMM*, vol. 32, no. 4, pp. 323–336, Aug. 2002.

[27] M. Charikar, K. Chen, and M. Farach-Colton, "Finding frequent items in data streams," in *EATCS ICALP*, 2002, pp. 693–703.

[28] M. Chen and S. Chen, "Counter tree: A scalable counter architecture for per-flow traffic measurement," in *IEEE ICNP*, 2015, pp. 111–122.

[29] T. Li, S. Chen, and Y. Ling, "Per-flow traffic measurement through randomized counter sharing," *Networking, IEEE/ACM Trans. on*, 2012.

[30] S. Cohen and Y. Matias, "Spectral bloom filters," in *ACM SIGMOD*, 2003, pp. 241–252.

[31] D. Ficara, A. D. Pietro, S. Giordano, G. Procissi, and F. Vitucci, "Enhancing counting bloom filters through huffman-coded multilayer structures." *IEEE/ACM Trans. Netw.*, vol. 18, pp. 1977–1987, 2010.

[32] R. Ben-Basat, G. Einziger, R. Friedman, and Y. Kassner, "Heavy hitters in streams and sliding windows," in *IEEE INFOCOM*, 2016.
[33] G. Einziger and R. Friedman, "Counting with tinytable: Every bit counts!" in *ACM ICDCN*, 2016.
[34] N. Hua, B. Lin, J. J. Xu, and H. C. Zhao, "Brick: A novel exact active statistics counter architecture," in *ACM/IEEE ANCS*, 2008, pp. 89–98.
[35] R. Morris, "Counting large numbers of events in small registers," *Commun. ACM*, vol. 21, no. 10, pp. 840–842, 1978.
[36] R. Stanojevic, "Small active counters," in *IEEE INFOCOM*, 2007, pp. 2153–2161.
[37] B.-Y. Choi, J. Park, and Z.-L. Zhang, "Adaptive random sampling for load change detection," *ACM SIGMETRICS*, pp. 272–273, 2002.
[38] C. Estan, K. Keys, D. Moore, and G. Varghese, "Building a better netflow," in *ACM SIGCOMM*, 2004.
[39] B. Manes, "Caffeine: A high performance caching library for java 8," *https://github.com/ben-manes/caffeine*.

APPENDIX

A. Theoretical Guarantees for Top-k Identification Using Randomized Admission Policy

To show the benefit of probabilistic admission, we describe a variant of RAP (see Section III) that aims to minimize the number of counters needed for top-k identification. We consider a setting in which the stream elements are i.i.d, i.e., each element is sampled independently and according to the same distribution. Also, we assume that elements are from a finite domain $\mathcal{U} = \{1, 2, \ldots, D\}$, and without loss of generality, the frequencies of the elements are

$$f_1 \geq f_2 \geq \ldots \geq f_k > f_{k+1} \geq f_{k+2} \geq \ldots \geq f_D.^2$$

For $r \in \{1, 2, \ldots, D\}$, we denote $F_r \triangleq \sum_{i=1}^{r} f_i$. This means that at each timestamp, element i will arrive with probability f_i and $\sum_{i=1}^{D} f_i = 1$. We note that this i.i.d. setting may not be applicable to certain streams exhibiting high time locality, such as packets going through a home router, but may resemble the traffic patterns appearing on major backbone routers. The goal of the top-k problem is then to identify the set of the most frequent elements $\{1, 2, \ldots, k\}$ with as few counters as possible. Our assumption is that the stream may be arbitrarily long, but we wish to guarantee that with probability 1 the algorithm will eventually identify all top-k items.

Formally, we say that algorithm A has successfully identified the top-k elements at time t if after the arrival of the t'th element, the k largest counters are allocated for items $\{1, 2, \ldots, k\}$. Since the actual items are random variables, we consider the probability, denoted $P^A_{m,k}(t)$, that algorithm A will successfully identify the top-k elements at time t if allocated with m counters. Finally, our benchmark would be the minimal number of counters that A requires to achieve

$$\lim_{t \to \infty} P^A_{m,k}(t) = 1.$$

We say that an item is a *tail item* if it is not amongst the top-k, i.e., is one of $\{k+1, \ldots, D\}$. We call the largest k counters the *main counters*, and the remaining ones *tail counters*. Our goal is then to ensure that after seeing infinitely many elements, the top-k elements will be guaranteed to be allocated with the main counters. For example, notice that D counters are enough for Space Saving, regardless of the actual frequencies $\{f_i\}$. However, the interesting case is when $m \ll D$.

We start by analyzing the number of counters Space Saving requires for this task. We compare to Space Saving, because it achieves asymptotic improvement over previous work [19] and, to the best of our knowledge, it is considered the state of the art. Notice that our analysis is somewhat different than the one presented in [19], as it further assumes that the stream is i.i.d., which reduces the number of counters required.

Algorithm 2 Space-Saving

Initialization: $C = \emptyset$

```
1: function ADD(Item x)
2:     if x ∈ C then
3:         Increment c_x
4:     else
5:         if |C| < M then
6:             c_x = 1
7:             C = C ∪ {x}
8:         else
9:             m = argmin_{y∈C} c_y
10:            C = C \ {m} ∪ {x}
11:            c_x ← c_m
12:            Increment c_x
13:        end if
14:    end if
15: end function
```

1) Conditions for Successful Space Saving Top-k Identification: Assume that we allocate a Space Saving instance with m counters, and would like to identity the top-k elements. In the algorithm, whose pseudo code appears in Algorithm 2, every arriving element is associated with a counter; if a counter was associated with the element prior to its arrival, the counter is incremented by 1; otherwise, the element "takes over" the minimal counter and increments it. Consequently, we are guaranteed that the top-k element will be associated with a counter upon arrival (unlike our algorithm described below). The only problem for Space Saving arises when a top-k element loses its main counter in favor of a "tail item", i.e., an item that is not within the top-k. Since k is the least-frequent element within the top-k, it is enough to consider whether it is guaranteed to have a counter. The key point of our analysis is observing that *if* all top-k are allocated with counters, they have a positive probability of forever being allocated with these counters if and only if the tail counters increase in a rate smaller than f_k. Observe that if the top-k elements reside within the main counters, the expected increment rate of the *sum* of tail counters is $\sum_{i=k+1}^{D} f_i = 1 - F_k$. This means that the average increment rate for a tail counter is $\frac{1-F_k}{m-k}$, and therefore the increment rate for the minimum among all tail counter is at most $\frac{1-F_k}{m-k}$. On the other hand, the counter allocated for item k increases with a rate of f_k. Next, we claim that having $f_k > \frac{1-F_k}{m-k}$ is a necessary and sufficient condition for Space Saving to identify the top-k. Illustration of the claim appears in Figure 7.

[2]If we do not wish to assume that f_k is strictly larger than f_{k+1}, we can replace our demand to finding a set of k items such that all of their frequencies are at least f_k. Such a model leads to similar results.

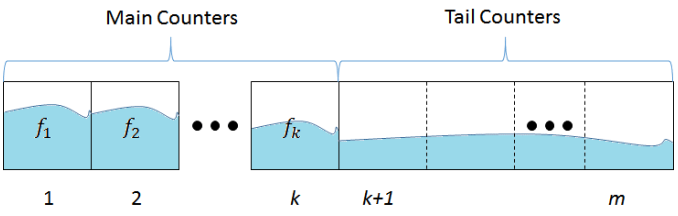

Fig. 7: The counters allocated for the top-k elements are increased with rates $f_1, \ldots, f_k$. The increase rate of the *minimal* tail counter is at most $\frac{1-F_k}{m-k}$.

Theorem A.1. *Space Saving successfully identifies the top-k using m counters if and only if $f_k > \frac{1-F_k}{m-k}$.*

Proof. We start by noting that with probability 1, if a top-k element is not allocated with a main counter, it will get such a counter in the future. This happens because we assumed $f_k > f_{k+1}$, which means that after a tail element takes over a main counter, the counter is incremented with a frequency of at most f_{k+1}, while the counter allocated with the missing top-k element (or the smallest counter if it is not currently allocated with one) increases at a rate of at least $f_k > f_{k+1}$.

We are left with showing that $f_k > \frac{1-F_k}{m-k}$ is a characterization of the ability of the top-k elements to seize the main counters without being evicted. As mentioned above, it is enough to argue that item k will be eventually allocated within the main counters to ensure all top-k items are successfully identified. Assume that k is allocated with a counter with value c_k, and let c_t be the value of the minimal tail counter. Notice that c_k is increased with rate f_k while c_t's rate is at most $\frac{1-F_k}{m-k}$. Consider the infinite Markov chain whose states represent the difference between c_k and c_t, i.e., state i represents the case of $c_k - c_t = i$. At any time an element arrives, it increments c_k with probability f_k, increments c_t with probability of at most $\frac{1-F_k}{m-k}$, and increments other counters otherwise. Therefore if we ignore other counters, the transition probabilities do not depend on the current state i and can be expressed as $\forall i : \Pr[i+1 \mid i] = \lambda \triangleq \frac{f_k}{\frac{1-F_k}{m-k}+f_k}$ and $\forall i : \Pr[i-1 \mid i] = 1 - \Pr[i+1 \mid i] = \mu \triangleq \frac{\frac{1-F_k}{m-k}}{\frac{1-F_k}{m-k}+f_k}$. Notice that

$$\Pr[i+1 \mid i] > 1/2 \iff f_k > \frac{1-F_k}{m-k}.$$

The stochastic process is illustrated in Figure 8.

Fig. 8: The probability of moving to a larger index state is λ and is not dependent on the process made so far or the current state index.

Since we know that top-k elements will obtain a main counter infinitely often, the question of successful identification narrows to the question "is there a positive probability that given a positive integer $n > 1$, such that if $c_k \geq c_t + n$, the process will *never* return to state 0?" (as then the main counter allocated for k may become minimal and lead to k being evicted). It is a known fact that a 1D random walk over the non-negative integers that goes left with probability μ and right with probability λ will return to 0 starting at state n is $\left(\frac{\mu}{\lambda}\right)^n$, which is strictly smaller than 1 for $\lambda > \mu$. As we are guaranteed to reach a positive difference infinitely often, when $\mu < 1/2$ we will eventually guarantee that item k will not be evicted if and only if

$$f_k > \frac{1-F_k}{m-k}. \tag{1}$$

This can further be expressed as a bound on the number of counters required by Space Saving:

$$m > k + \frac{1-F_k}{f_k}. \tag{2}$$

□

2) Conditions for Successful RAP' Top-k Identification: In this subsection, we present a variant of the Randomized Admission Policy algorithm (see Algorithm 1), called RAP', that aims to minimize the number of counters required for top-k identification. RAP' takes advantage of the fact that we can "slow" the frequency in which the smallest counter is incremented for achieving better identification results. Formally, RAP' acts similarly to RAP, except that upon arrival of a non-monitored element, the probability in which it will be allocated with the minimal counter is a constant P, and do not depend on the value of the minimal counter.[3]

Notice that these probabilistic allocations has both positive and negative effects on the possibility of successful identification. On the positive side, since the minimal counter is incremented slower than in Space Saving, we can relax the $f_k > \frac{1-F_k}{m}$ constraint a bit. Alas, RAP' is not guaranteed that a top-k element will obtain a counter infinitely often. This means that for RAP' to successfully identify the top-k we need to impose two constraints:

1) A top-k item which is not currently allocated with a counter will get one (w.p. 1).
   Notice that if a top-k element is not allocated with a counter (recall that now we have $m-1$ allocated counters), then the minimal, non-allocated counter increases with rate of at least $P \cdot (f_k + 1 - F_m)$. In contrast, the slowest allocated counter rate cannot exceed f_m. Thus, the imposed constraint is

$$f_m < P \cdot (f_k + 1 - F_m). \tag{3}$$

   This gives us a lower bound on the value of P:

$$P > \frac{f_m}{f_k + 1 - F_m}. \tag{4}$$

[3]If there exists more than a single minimum-value counter, then the arriving element is admitted with probability 1.

For the sake of simplifying the calculations, we impose a stronger restriction on the admission probability and consider P to satisfy

$$P > \frac{f_m}{f_k}. \tag{5}$$

2) When all top-k elements are allocated with counters, there exists a positive probability that they will never be evicted from the table. This is the front in which RAP' has advantage over Space Saving. Since the minimal counter is only incremented with sampling rate P, the rate in which the smallest counter within the tail increases is at most

$$\frac{\sum_{i=k+1}^{m} f_i + P \cdot \sum_{i=m+1}^{D} f_i}{m-k} = \frac{F_m - F_k + P \cdot (1 - F_m)}{m-k}.$$

Meanwhile, the counter associated with item k is incremented in rate of f_k. This means that our constraint is:

$$f_k > \frac{F_m - F_k + P \cdot (1 - F_m)}{m-k} \tag{6}$$

Notice that for $P = 1$, Inequality (5) trivially holds (and thus Space Saving is guaranteed to get counters allocated for top-k elements infinitely often), while Inequality 6 degenerates into Inequality (1). In the following subsection, we analyze how a smart choice of the increment probability P reduces the required number of counters.

3) Performance Comparison for Zipf Distributed Streams: In many of the previous works, Zipf distributed streams served as a popular benchmark for algorithms comparison due to its nice mathematical properties. Here, we continue this line and compare the performance of Space Saving and RAP' on i.i.d Zipf distributed streams with varying skews. We start with a formal definition of a Zipf stream.

Definition Denote $\Gamma_\alpha(D) = \sum_{i=1}^{D} i^{-\alpha}$. A stream will be called an i.i.d Zipf stream with skew α over domain D if all of its elements are sampled independently and follow the distribution in which item $i \in \{1, 2, \ldots, D\}$ appears with probability $f_i = \frac{i^{-\alpha}}{\Gamma_\alpha(D)}$.

Traditionally, streams with skew $\alpha \in (0, 1]$ are called "mildly skewed" or "heavy tailed", while larger skews grants the streams the title "highly skewed". Streams in which every item is selected with uniform probability (skew=0) are called "uniform". These names differentiate highly skewed streams, in which a small number of elements consists most of the stream, and heavy tailed ones, where most of the arriving elements are tail items. This property is also observable from the behavior of $\Gamma_\alpha(D)$; for $\alpha > 1$, $\Gamma_\alpha(D)$ convergence to a constant as D grows (e.g., $\Gamma_2(\infty) \approx 1.645$); for $\alpha = 1$, $\Gamma_1(D) \approx \ln(1.78D)$; lastly, for $\alpha < 1$, we have $\Gamma_\alpha(D) = \frac{D^{1-\alpha}}{1-\alpha} + O(1)$.

For heavily skewed streams, Space Saving is known to be optimal [19]. For more mildly skewed streams, we show that RAP' could asymptotically improve the number of counters required for identifying the top-k elements. This also provides theoretical grounds to the poor empirical performance of Space Saving when evaluated on heavy tailed workloads in Section V.

Assuming a Zipf distributed stream, the condition for Space Saving to converge, as appears in (2), then becomes:

$$\begin{aligned} m &> k + \frac{1 - F_k}{f_k} = k + \frac{1 - \frac{\Gamma_\alpha(k)}{\Gamma_\alpha(D)}}{\frac{k^{-\alpha}}{\Gamma_\alpha(D)}} \\ \implies m &> k + k^\alpha \left(\Gamma_\alpha(D) - \Gamma_\alpha(k)\right) \end{aligned} \tag{7}$$

For analyzing RAP's performance, we will select the value of P based on the skewness of the stream, as discussed below. When plugging the Zipf distribution into the first RAP' constraint (see (5)), we get

$$\begin{aligned} P &> \frac{f_m}{f_k} = \left(\frac{m}{k}\right)^{-\alpha} \\ \implies m &> k \cdot P^{-\frac{1}{\alpha}}. \end{aligned} \tag{8}$$

Similarly, the second constraint (see (6)) is now:

$$\begin{aligned} f_k &> \frac{F_m - F_k + P \cdot (1 - F_m)}{m-k} \\ \iff m &> k + k^\alpha \Big(\Gamma_\alpha(m) - \Gamma_\alpha(k) \\ &\quad + P \cdot (\Gamma_\alpha(D) - \Gamma_\alpha(m)) \Big) \end{aligned}$$

In order to simplify the right hand side of the inequality, we impose a stronger bound on m and require is to satisfy:

$$m > k + k^\alpha \cdot (\Gamma_\alpha(m) - \Gamma_\alpha(k) + P \cdot \Gamma_\alpha(D)) \tag{9}$$

4) Heavy Tailed Streams: In this section, we assume that $\alpha \in (0, 1)$ is fixed and that $k = o(D^{\frac{\alpha}{1+\alpha}})$ and analyze the number of counters required for Space Saving and RAP' for successfully identifying the top-k items.[4] We start by using the explicit formula of $\Gamma_\alpha(\cdot)$ for (7):

$$\begin{aligned} m_{SS} &= k + k^\alpha \left(\Gamma_\alpha(D) - \Gamma_\alpha(k)\right) \\ &= k + k^\alpha \frac{D^{1-\alpha} - k^{1-\alpha} + \Theta(1)}{1-\alpha} \\ &= \frac{k^\alpha \cdot D^{1-\alpha} - \alpha \cdot k + \Theta(k^\alpha)}{1-\alpha} = \Theta(D^{1-\alpha}) \end{aligned}$$

Thus, we established that the number of counters required for Space Saving is $m_{SS} = \Omega(D^{1-\alpha})$.

For RAP', we choose the admission probability to be

$$P \triangleq D^{\frac{\alpha^2 - \alpha}{1+\alpha}}. \tag{10}$$

Notice that $P \in (0, 1)$ is a valid probability. Next, we will show that using

$$m_{RAP'} \triangleq c \cdot k \cdot D^{\frac{1-\alpha}{1+\alpha}} \tag{11}$$

counters, where c is a (large enough) constant, we can satisfy both (8) and (9), thus $m_{RAP'}$ counters are enough for successful identification of the top-k elements.

[4] In practice, values of k are typically very small and may be considered sub-polynomial in D.

Constraint (8) requires that $m > k \cdot P^{-\frac{1}{\alpha}}$. Plugging in (10) and (11), we get:

$$m_{RAP'} = ck \cdot D^{\frac{1-\alpha}{1+\alpha}} > k \cdot P^{-\frac{1}{\alpha}},$$

as required. Next, we show an inequality that will be useful later:

$$\begin{aligned}(m_{RAP'})^{1-\alpha} &= (ck \cdot D^{\frac{1-\alpha}{1+\alpha}})^{1-\alpha} \\ &= (D^{\frac{1}{1+\alpha}} \cdot \frac{ck}{D^{\frac{\alpha}{1+\alpha}}})^{1-\alpha} < D^{\frac{1-\alpha}{1+\alpha}}, \end{aligned} \tag{12}$$

where the last inequality holds for large enough c.

Finally, we show that our choice of P and $m_{RAP'}$ also satisfies (6):

$$\begin{aligned}&k + k^{\alpha} \cdot (\Gamma_{\alpha}(m) - \Gamma_{\alpha}(k) + P \cdot \Gamma_{\alpha}(D)) \\ &= k + \frac{k^{\alpha} \cdot \left(m^{1-\alpha} - k^{1-\alpha} + D^{\frac{\alpha^2-\alpha}{1+\alpha}} \cdot D^{1-\alpha} + \Theta(1) \right)}{1-\alpha} \\ &= \frac{k^{\alpha} \cdot \left(m^{1-\alpha} + D^{\frac{1-\alpha}{1+\alpha}} \right) - \alpha k + \Theta(k^{\alpha})}{1-\alpha} \\ &< \frac{2k^{\alpha} \cdot D^{\frac{1-\alpha}{1+\alpha}} - \alpha k + \Theta(k^{\alpha})}{1-\alpha} \\ &= m_{RAP'} \cdot \frac{2}{c(1-\alpha)} + O(k) < m_{RAP'} \end{aligned} \tag{13}$$

Since we have shown that our selection of admission probability and number of counters satisfies both constraints, we have proved that RAP' requires only $O(D^{\frac{1-\alpha}{1+\alpha}})$ counters to successfully identify the top-k hitters on a Zipf stream with skew α.

We conclude that for the problem of identifying top-k over i.i.d. heavy tailed streams, Space Saving requires $\Theta(D^{1-\alpha})$ while RAP' requires $\Theta(D^{\frac{1-\alpha}{1+\alpha}})$. Notice that for values of α that are close to 1, this is nearly a quadratic space reduction. For example, consider trying to find the top-32 flows on a backbone router whose traffic is approximately Zipf0.8 with domain of $D = 2^{64}$ elements. Space Saving requires about $570K$ counters; in contrast, RAP' could allocate roughly $44K$ counters to achieve the same. The admission probability for these input parameters is slightly less than 2%.

5) Skew=1 Streams: Heavy tailed streams are usually not analyzed in the literature, perhaps because the existing algorithm cannot find the top-k elements in these using a reasonable amount of counters (see Section V). However, skew 1 Zipf streams were analyzed in some previous works for both Space Saving and Count Sketch [19], [27]. In this section, we show that by introducing an admission probability, RAP' is able to achieve asymptotic space improvement on i.i.d. Zipf streams. The convergence condition for Space Saving, which appears in (7) is now:

$$\begin{aligned}m &> k + k\left(\Gamma_1(D) - \Gamma_1(k)\right) \\ &\approx k + k \ln\left(1.78\frac{D}{k}\right) = \Theta(\log D). \end{aligned} \tag{14}$$

For RAP', we choose the admission probability to be

$$P \triangleq \sqrt{\frac{1}{\ln D}}. \tag{15}$$

We show that using probabilistic admission, we reduce the number of required counters to $m_{RAP'} \triangleq c \cdot k\sqrt{\ln D} = O(\sqrt{\log D})$, where c is a positive constant. We start by showing that this choice of $m_{RAP'}$ and P satisfies (8):

$$m_{RAP'} = c \cdot k\sqrt{\ln D} > k \cdot P^{-1} \tag{16}$$

Next, we consider (9):

$$\begin{aligned}&k + k \cdot (\Gamma_1(m) - \Gamma_1(k) + P \cdot \Gamma_1(D)) \\ \approx &k + k \cdot \left(\ln\left(1.78\frac{m}{k}\right) + \sqrt{\frac{1}{\ln D}} \ln(1.78D) \right) \\ = &k + k \cdot \left(\ln\left(c \cdot \sqrt{\ln D}\right) + \sqrt{\ln D} + O(1) \right) < m_{RAP'}, \end{aligned}$$

where the last inequality holds for large enough c.

We conclude that by introducing a probabilistic admission filter, one can reduce the number of counters required for successful top-k identification over skew=1 Zipf streams from $O(\log D)$ to $O(\sqrt{\log D})$.

6) Discussion and Future Work: In this section, we have shown how the concept of admission filters, introduced in Section III can be adapted for reducing the number of counters needed for top-k identification. Nevertheless, our analysis has a few drawbacks; first, we are only able to provide theoretical analysis for i.i.d. streams, which may not truly represent the all practical settings; second, our analysis assumes that the stream may be arbitrarily long, and only considers *eventual convergence*. While this may fit massive data streams, in cases where we wish to process smaller streams, perhaps because we reset the process every once in a while to provide freshness, our analysis does not hold. Lastly, we have assumed a prior knowledge of the data skew. In practice, one can estimate the skew from the data, but this will require the admission probability to be adaptive.

We plan to evaluate RAP' on real data traces and compare it with existing techniques, and specifically with RAP (see Algorithm 1) whose admission probability was optimized for frequency estimation but proved effective also for top-k. We also wish to find a method for dynamically adapting the admission probability without assuming knowledge about the stream length or skew.

B. Limited Associativity Impact

In this appendix, we compare the performance of d-Way RAP for different values of d. We evaluate the associativity levels effect over several metrics which are presented in detail in Section V-C and Section V-D. These include the following:

1) On-Arrival Mean Square Error, in which every arriving element is queried and we compute the average square error.
2) The percentage (recall) of elements within the top-32 successfully using various space allocations.

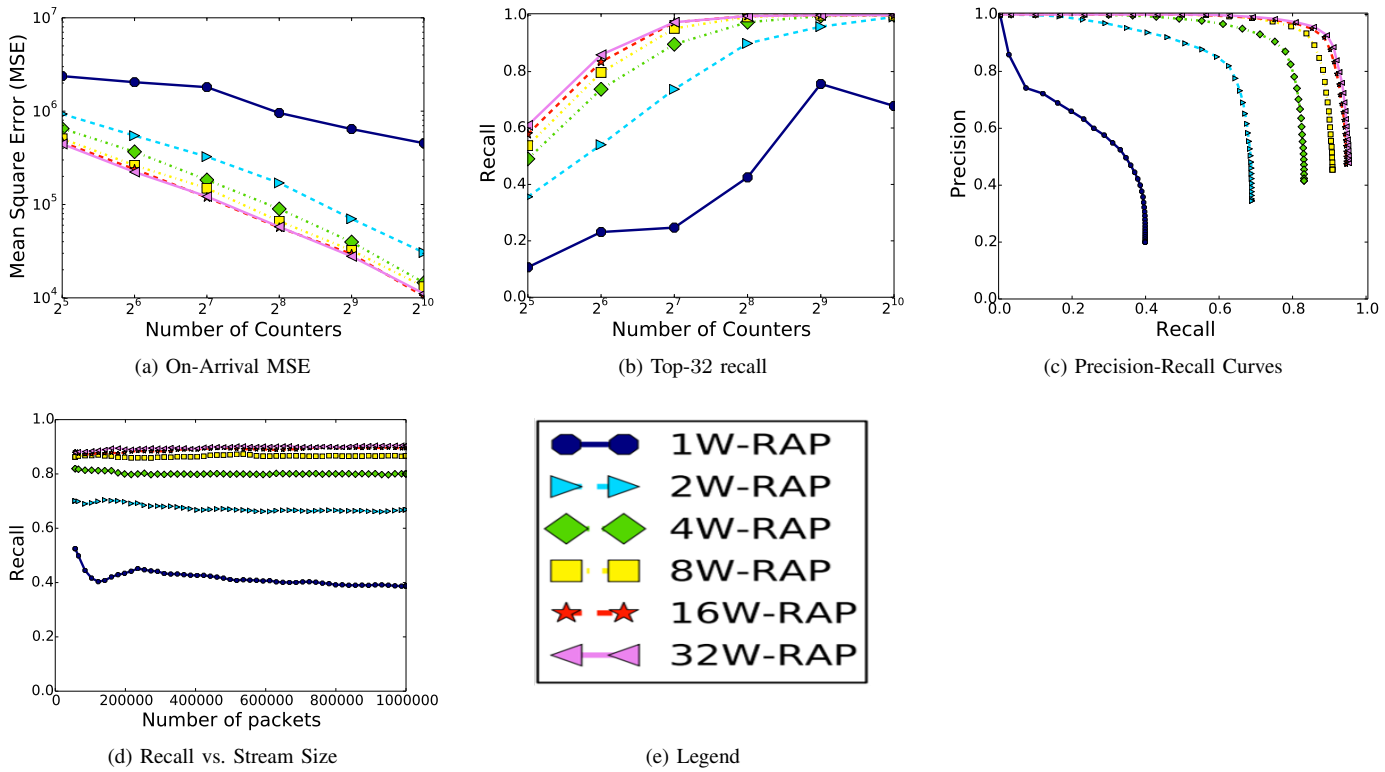

(a) On-Arrival MSE

(b) Top-32 recall

(c) Precision-Recall Curves

(d) Recall vs. Stream Size

(e) Legend

Fig. 9: Comparison of the performance of d-Way RAP for different associativity levels.

3) The recall for identifying the top-512 elements using 1024 counters, compared with the number of observed packets.
4) The precision-recall curve for identifying the top-512 elements using 1024 counters.

Figure 9 shows the performance of the different associativity levels, averaging over 10 batches of 1M packets each from the CAIDA [21] dataset. The results show a diminishing return pattern as associativity is increased; while 1W-RAP performs rather poorly, 2W-RAP is already comparable with the previous algorithms, 4W and 8W offer increased accuracy while 16W-RAP works almost as good as the $32W$. Further, our evaluation in Section V-C and Section V-D shows that $16W$-RAP is roughly comparable to the fully associative RAP. Our experiments suggest that associativity of 16 counters per set is a highly attractive alternative to complete associativity as it does not require any sophisticated data structures (as suggested in [18], [19], [32]). By not using data structures, we get both simpler implementation, as well as reduction in the memory overhead they require, which can be used for allocating the algorithms with additional counters for increased accuracy.